An Algorithm for the Reduction of Aerosol Model Noise in SeaWiFS Data

Bryan A. Franz

1 August 2001

Updated: 1 June 2002

I. Introduction and Recommendation

The purpose of this analysis is to investigate methods to reduce small-scale (pixel-to-pixel) spatial variability in the SeaWiFS Level-2 derived products. These effects appear to be the result of noise in the residual aerosol signal of the NIR bands, which gives rise to excessive aerosol model switching. The goal is to develop a procedure to smooth the unrealistic variability in aerosol type retrievals, without significantly altering the retrieved aerosol optical thickness in the NIR.

Since this data smoothing algorithm will be applied on the entire SeaWiFS mission archive, the method must be efficient in both processing time and memory usage. The algorithm will be operated within the Multi-Sensor Level-1 to Level-2 software (MSL12). MSL12 includes a capability for buffering of multiple lines of input TOA radiance data with associated flags, path geometries, and pre-computable atmospheric quantities (e.g., cloud flags, Rayleigh path radiances), and it provides a mechanism for data filtering which allows for the replacement of TOA radiances in each pixel of a scan line with new values derived through evaluation of a user-specified filter function over an m-pixel by n-line sliding window.

It is recommended that we apply a 5x5 diamond filter to the Rayleigh-subtracted TOA radiances in the 765 nm band, targeted to minimize the spatial variability of atmospheric epsilon. A description of this concept and the basis for recommendation is provided in the sections which follow.

II. Characterization of the Problem

The small-scale variability in SeaWiFS aerosol type retrievals is clearly illustrated by the spatial distribution of Ångstrom exponents in Figure 1. This is a 1-km resolution image extracted from the Western Mediteranean on 28 February 1998, and expanded 3-to-1 to show the individual pixels. The north coast of Algeria appears in the lower right corner. The dynamic range of the Ångstrom exponents in this image is 0.0 to 1.5, which covers almost the entire range of aerosol types represented in the SeaWiFS aerosol model suite, as shown in Table 1. It is reasonable to assume that the pixel-to-pixel variability from low-humidity tropospheric aerosols (red pixels) to coastal aerosols (green pixels) is physically unrealistic. There is also some indication of horizontal striping, which may be an effect of residual mirror-side reflectance differences.

Since SeaWiFS is an ocean color mission and not an atmospheric sciences mission, we can accept physically unrealistic variability in the aerosol type if that variability does not

significantly impact the retrieval of ocean optical properties. Figure 2 shows the retrieved normalized water-leaving radiance at 443 nm, for the same subscene as Figure 1. The dynamic range of the image has been stretched over the narrow range from 0.4 to 1.0 mW cm^{-2} μm^{-1} sr^{-1}, to better show the small-scale variability. The image shows frequent speckling (e.g., green pixels in the blue gyre), as well as some suggestion of horizonal striping. These artifacts are similar to what we see in the Ångstrom image.

To verify that the artifacts in the nLw(443) image are not simply due to instrumental noise in the 443 nm channel, a simple test can be performed. The Ångstrom coefficients shown in Figure 1 are directly computed by MSL12 from the selected aerosol model or model pair. The average Ångstrom coefficient in our test subscene is 0.90, which falls between the C50 and T99 models listed in Table 1. A histogram of the selected aerosol models also shows that the C50 is the most common lower-bounding aerosol model selected, and T99 is the most common upper bounding model. So, to remove the effect of aerosol model selection noise while retaining roughly the same aerosol radiance contribution, we can recompute the nLw(443) image of Figure 2 using a fixed aerosol model consisting of equal parts C50 and T99. The result is shown in Figure 3.

The speckling noise and most of the striping observed in Figure 2 has been effectively eliminated in Figure 3. Of courese, fixing the aerosol type is not a solution to the problem, but the demonstration serves to isolate the source of the noise. The aerosol type is determined by the two NIR channels of SeaWiFS, at 765 and 865 nm. In effect, the 865 nm channel controls the aerosol concentration, and the ratio of the 765 to 865 nm channels controls the aerosol model selection. By fixing the aerosol model, we have eliminated the influence of the 765 nm channel, but we have retained the influence the 865 nm channel which determines aerosol concentration and thus directly scales aerosol path radiance and aerosol optical thickness. We can conclude from this that the speckling seen in the nLw(443) image is primarily caused by relative variability between the two NIR channels.

Table 1: *SeaWiFS standard aerosold model suite for reprocessing #3.*

Model #	Identifier	Long Name	Ångstrom(510,865)
1	O99	Oceanic 99% Relative Humidity	-0.0867
2	M50	Maritime 50% Relative Humidity	0.5032
3	M70	Maritime 70% Relative Humidity	0.4019
4	M90	Maritime 90% Relative Humidity	0.2054
5	M99	Maritime 99% Relative Humidity	0.0904
6	C50	Coastal 50% Relative Humidity	0.7607
7	C70	Coastal 70% Relative Humidity	0.6456
8	C90	Coastal 90% Relative Humidity	0.4046
9	C99	Coastal 99% Relative Humidity	0.2251
10	T50	Tropospheric 50% Relative Humidity	1.5294
11	T90	Tropospheric 90% Relative Humidity	1.3862
12	T99	Tropospheric 99% Relative Humidity	1.1936

The relative variability between the NIR channels is further illustrated by Figure 4, which shows the ratio of TOA radiances at 765 and 865 nm, Lt(765)/Lt(865). This ratio would be expected to show large-scale trends due to spectral differences in the Rayleigh and average aerosol phase functions, as well as possible meso-scale variability due to changes in aerosol type. The small-scale speckling and line-to-line striping are clearly not physical.

The striping effect can be explained by spectral differences in mirror side reflectivity. The scanning mechanism of the SeaWiFS instrument make use of a double-sided mirror, with adjacent lines of data collected from opposite sides of that mirror. If, for a given wavlength, the reflectivity differes between the two mirror sides, striping will occur. The standard SeaWiFS calibration includes a correction for this reflectivity difference as a function of both wavelength and time. The residual mirror-side effects that we can see in Figure 4 are due to differences at the one-count level, below the level of correctibility for a single measurement of a digital instrument.

In terms of radiance units, the value of one count is a function of many factors (e.g.: time, focal-plane temperature, scan modulation). In the test scene of Figures 1-4, one count is equivalent to about 0.5% of the mean radiance in either NIR channel, which translates to as much as 1% in the TOA radiance ratio. So, the digitization error is small but not insignificant. In fact, a 1% difference in the TOA radiance ratio represents more than 10% of the dynamic range of Figure 4. The problem is further exacerbated within the atmospheric correction process, since the weight of a single count will increase as we subtract the contributions from non-aerosol sources such as the Rayleigh path radiance. For our test scene, the Rayleigh component contributes more than 70% of the TOA radiance in the NIR, so a 1-count difference translates to more than 2% of the Rayleigh-subtracted radiance at 765 or 865 nm. These numbers are summarized in Table 2.

Table 2: *Relative value of one count on median TOA radiance and median Rayleigh-subtracted TOA radiance for the test subscene of Figures 1-4.*

Band nm	Radiance/cnt G^i	Lt	Lr	Lt - Lr	100*G/Lt %	100*G/(Lt-Lr) %
412	0.0134	7.6044	7.1116	0.4953	0.1762	2.7051
443	0.0129	6.7128	6.0407	0.6403	0.1924	2.0174
490	0.0099	4.8361	4.1749	0.6097	0.2039	1.6172
510	0.0087	4.0213	3.5053	0.4615	0.2165	1.8869
555	0.0071	2.7150	2.4673	0.2083	0.2625	3.4218
670	0.0039	1.1183	0.9597	0.1461	0.3469	2.6560
765	0.0026	0.5211	0.4038	0.1122	0.5077	2.3582
865	0.0020	0.3237	0.2320	0.0898	0.6329	2.2803

i) mW cm^{-2} μm^{-1} sr^{-1} cnt^{-1}, median value specific to test subscene.

Figure 1: *Ångstrom Coefficient at 510 nm, standard SeaWiFS reprocessing #3.*

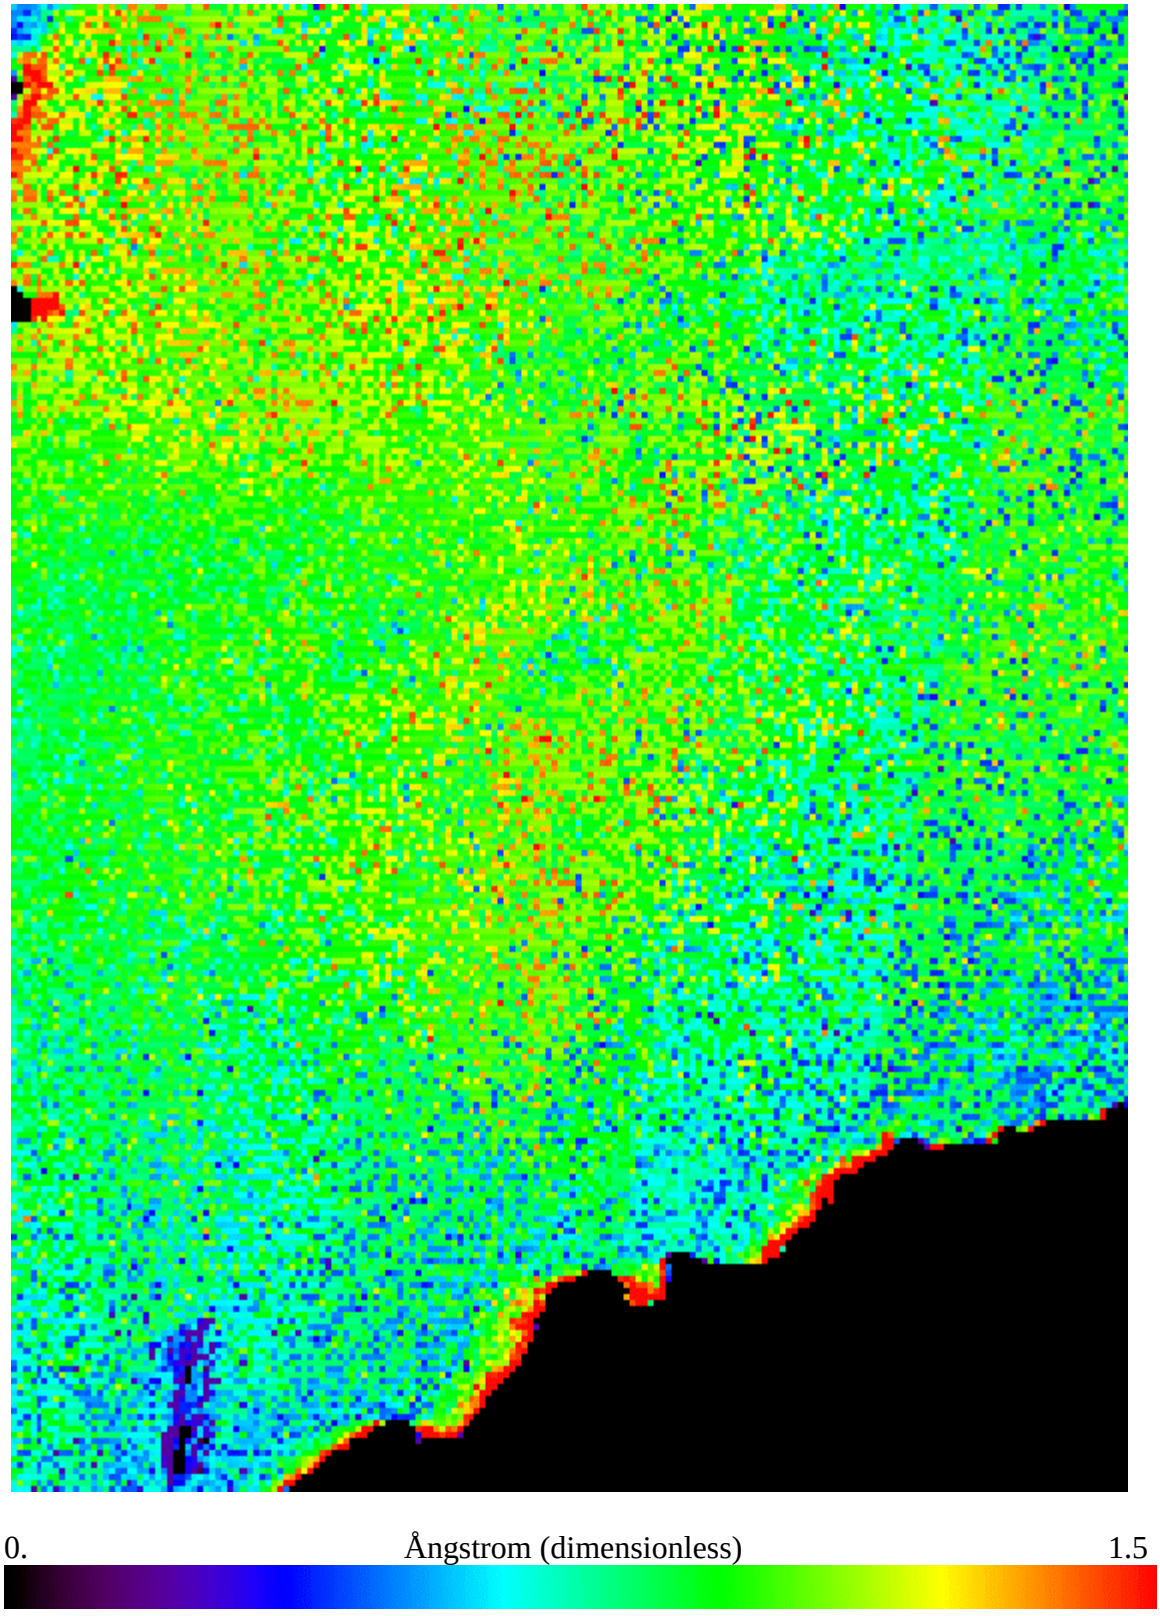

Figure 2: *Normalized water-leaving radiance at 443 nm, standard SeaWiFS processing .*

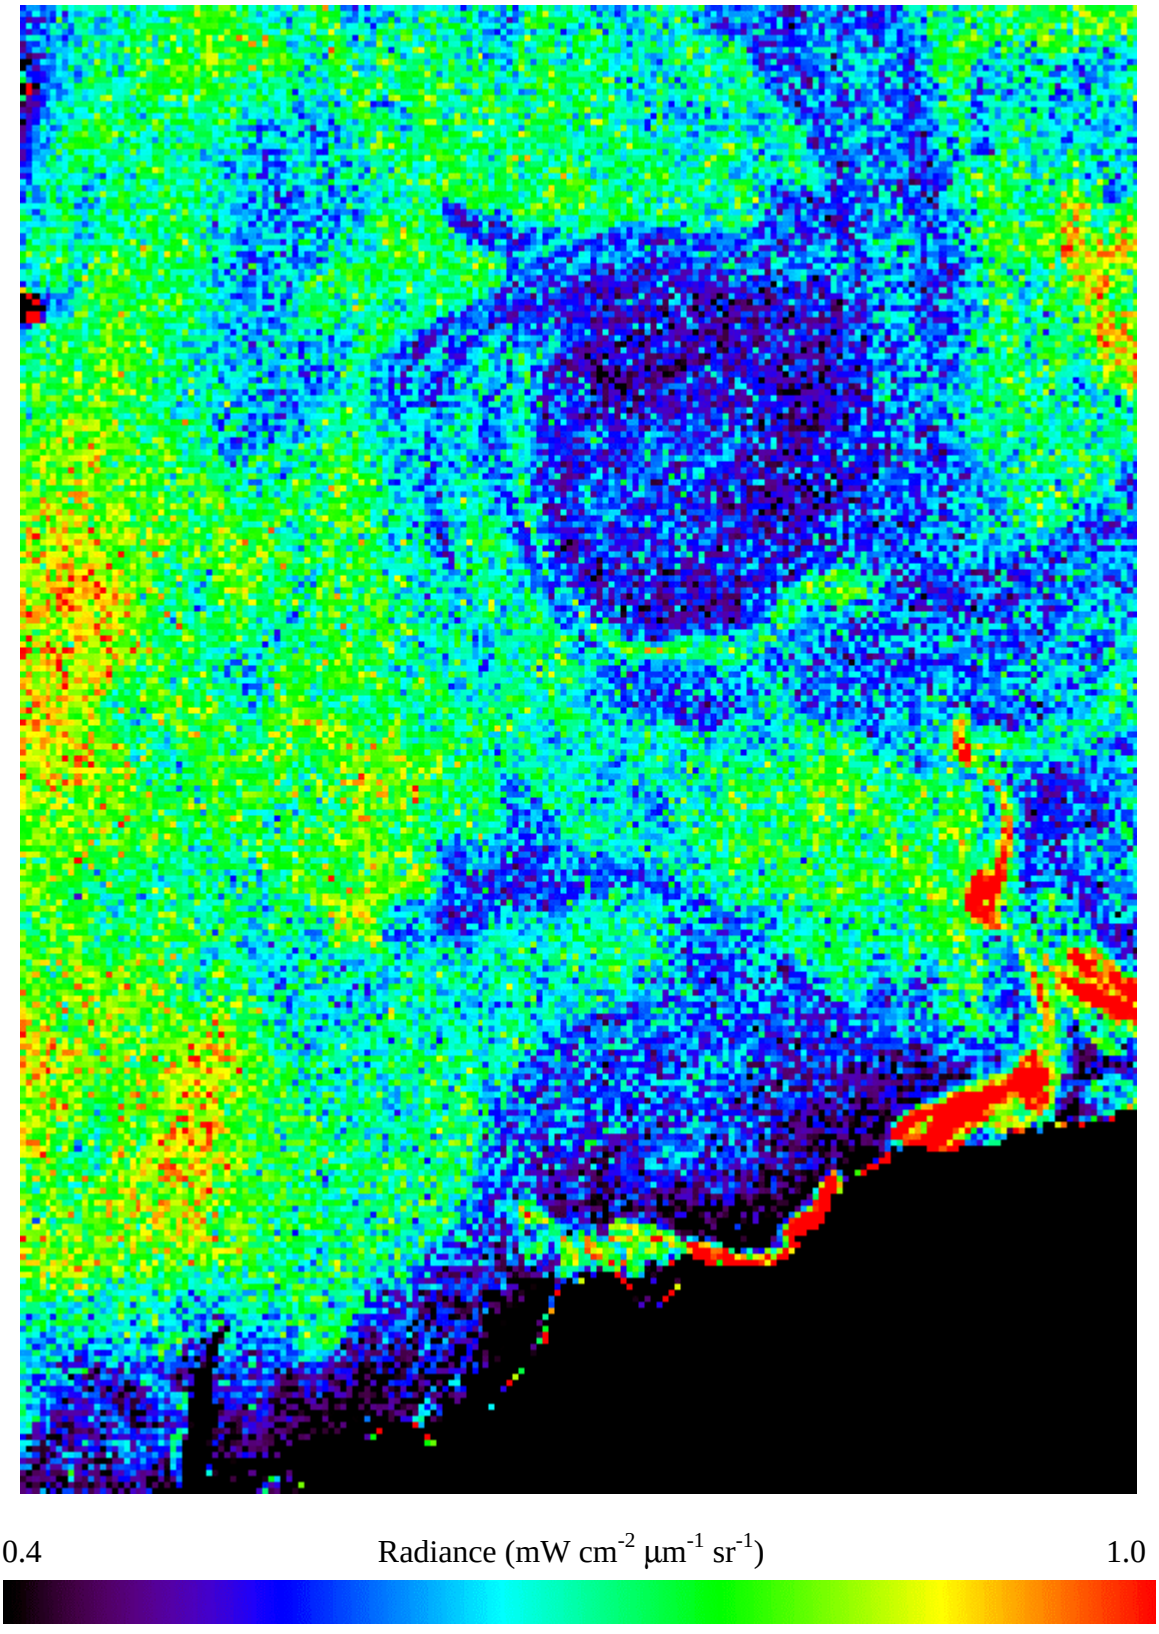

Figure 3: *Normalized water-leaving radiance at 443 nm, fixed aerosol models* .

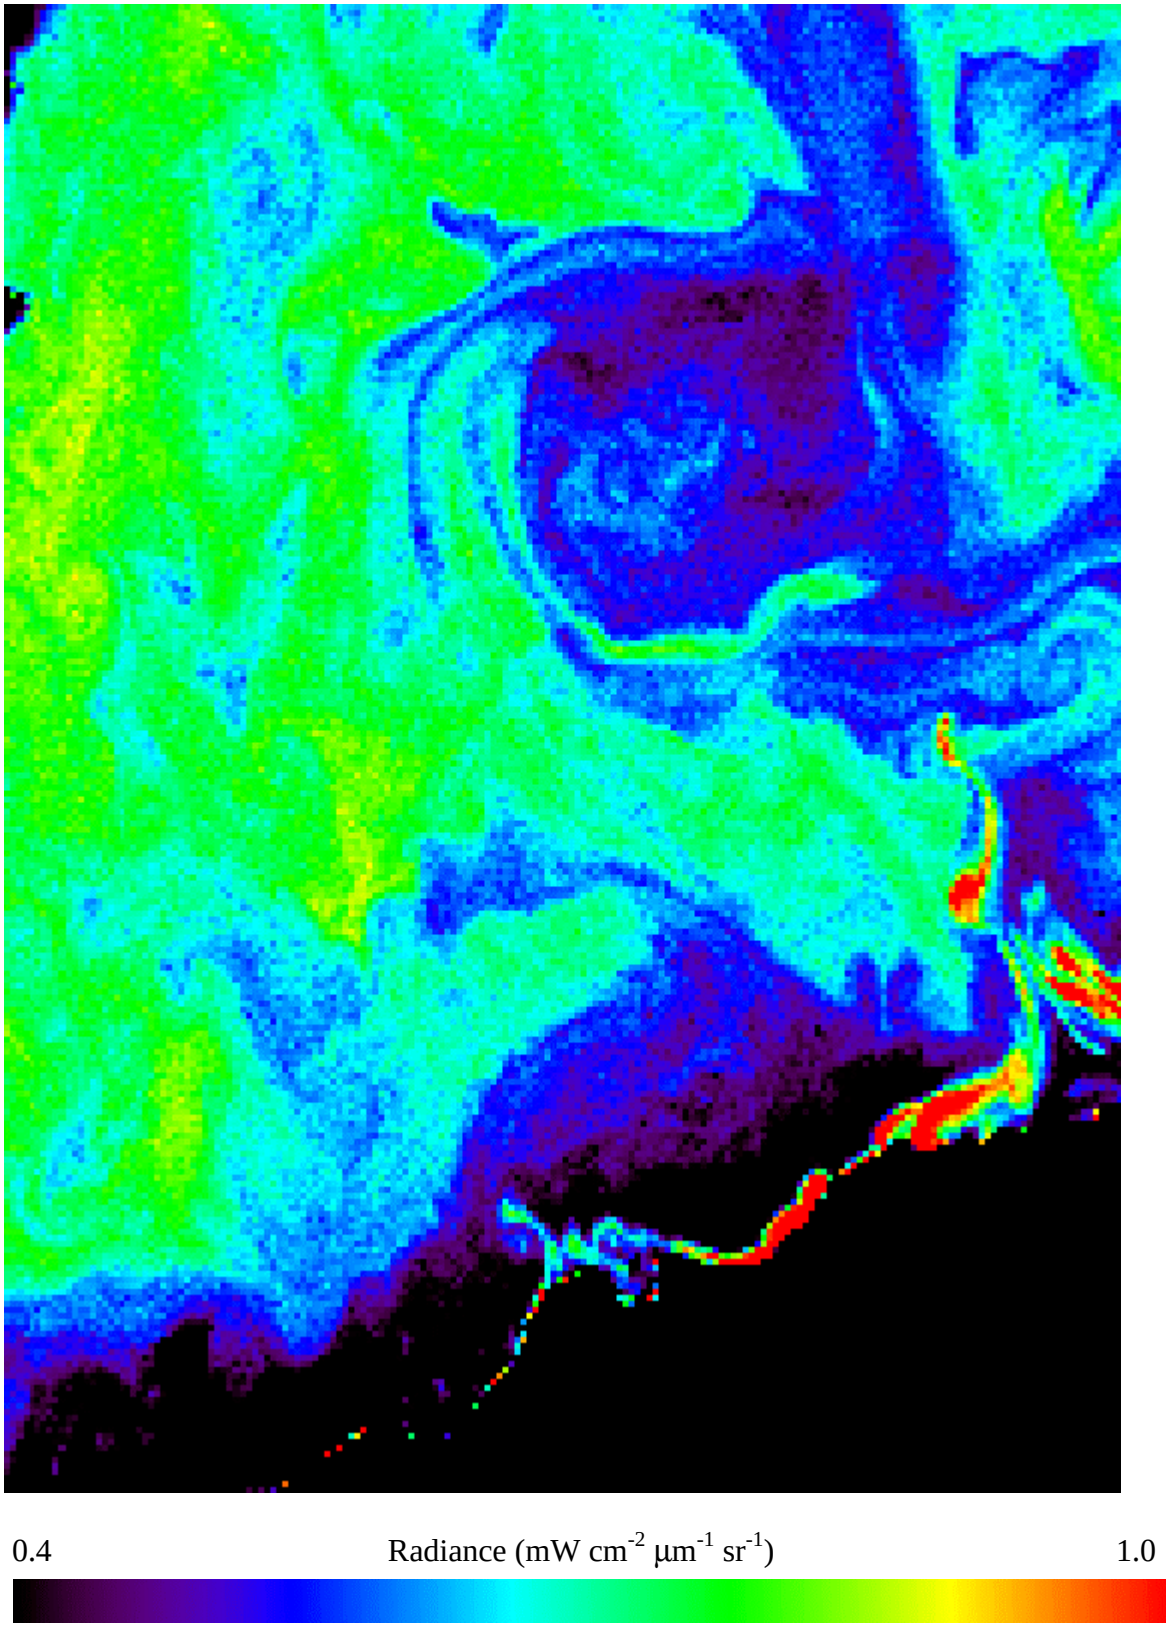

Figure 4: *Ratio of TOA radiance, Lt(765)/Lt(865) .*

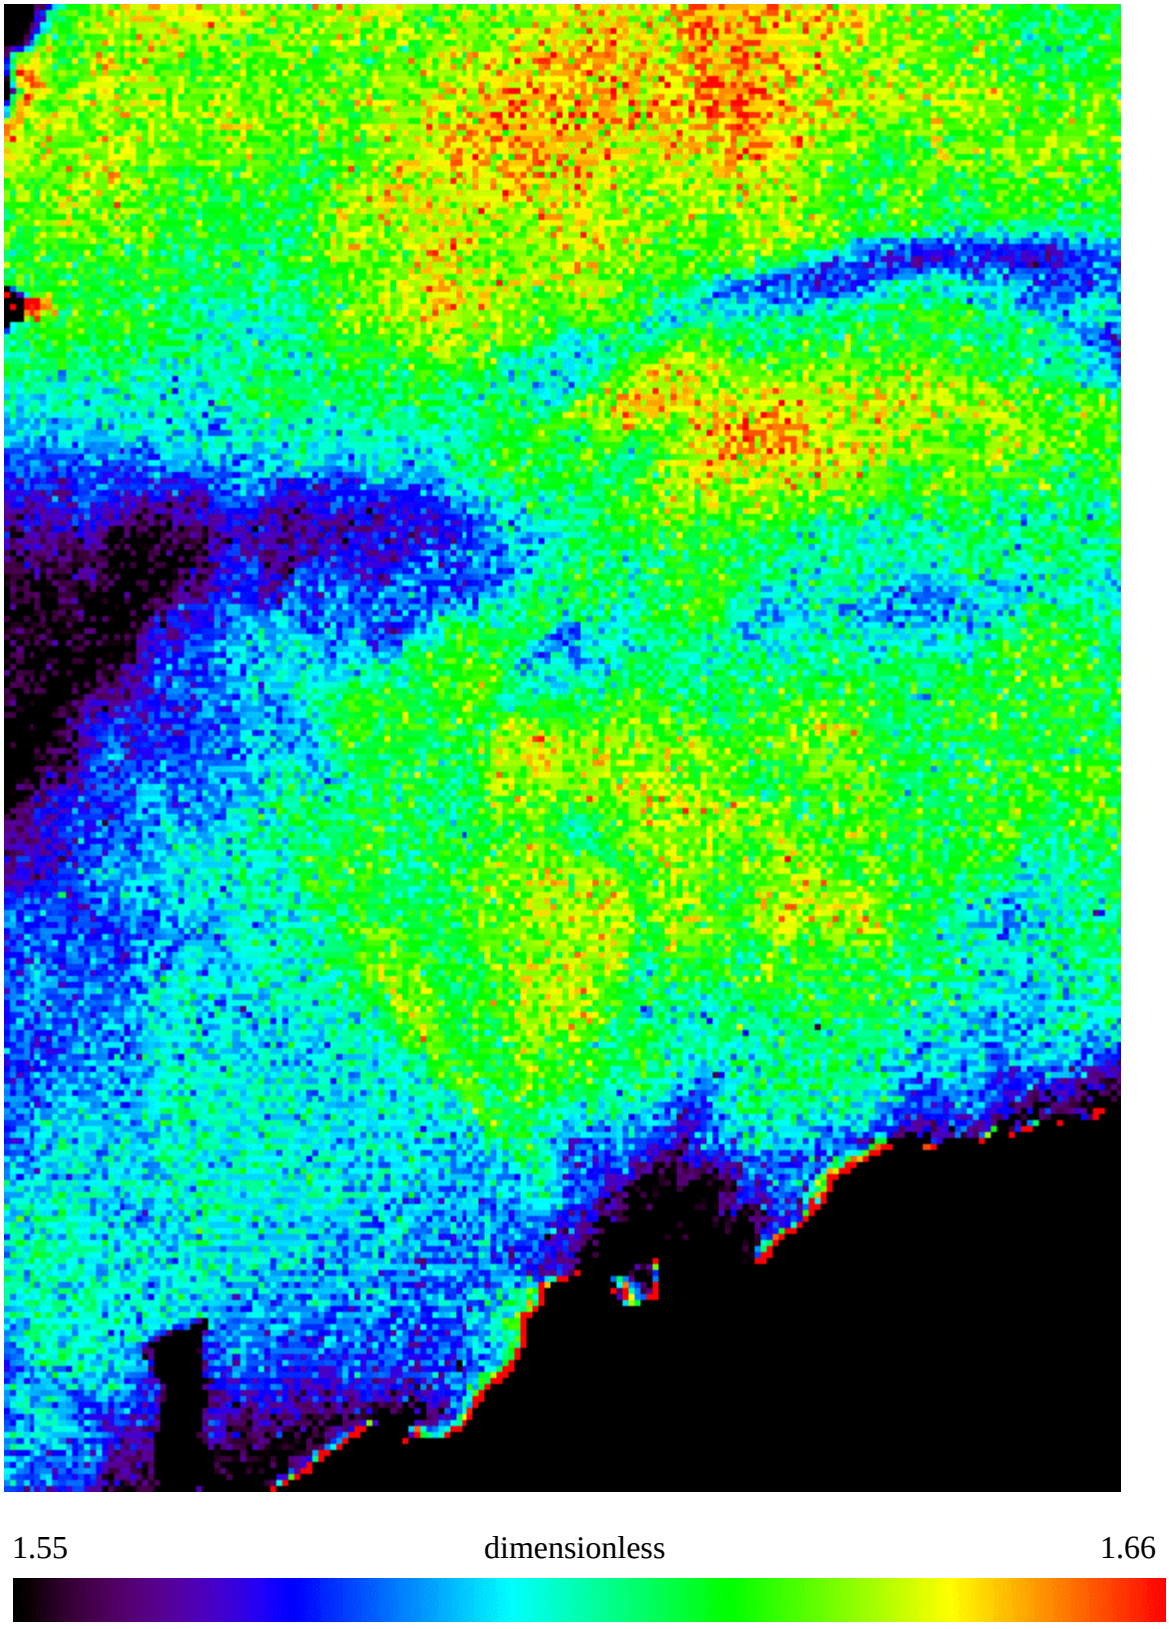

III. Filtering Algorithm Development

From the previous discussion it is evident that a primary source of noise in the retrieved ocean optical properties can be directly associated with the relative variability between the two NIR channels. The obvious solution to this problem is to apply some variety of spatial filter to the data, to remove outliers and/or increase signal-to-noise (S/N). The filtering mechanism within MSL12 provides an m-pixel by n-line sliding window, within which we can make any variety of computations and replace the central pixel TOA radiance or other pre-computed quantities (e.g., flags) as needed to ensure a smooth field or to eliminate highly deviant values. For example, we can use this framework to replace the TOA radiance in each pixel of each band with the median or mean TOA radiance over an m x n window, centered on that pixel.

If the speckling in the aerosol type is due to the occurence of radiance measurements which significantly deviate from surrounding measurements, then a median filter would be most effective at removing such outliers. However, if the problem is better characterized by digitization noise or insufficient signal-to-noise, then spatial averaging should improve the results. We know that the striping artifacts are due to digitization noise, so some averaging will be neccessary to reduce that effect. A simple mean spatial filter on the two NIR channels would serve to increase the signal-to-noise in each channel, and reduce the relative vaiability between the two channels.

Unfortunately, averaging of the 865 nm channel can create other problems. As stated previously, this band effectively controls the aerosol concentration. It is possible for the standard SeaWiFS aerosol correction process to mis-identify sub-pixel clouds, thin cirrus, fog, or even residual white-caps as thick aerosols. The extraneous radiance information is captured in the 865 nm channel, and extrapolated to the other bands through an aerosol model. The resulting aerosol radiance contribution derived for the visible channels may not be a perfect representation for such mis-identified contributions, but it should be reasonable because all such sources are relatively white in color. However, if we spatially average the 865 nm radiance, we will spread the erroneous aerosol contribution over neighboring pixels, which can cause significant over-subtration of the aerosol contribution in the visible channels. One solution to that problem would be to spatially average all bands equally, but that means we will reduce the resolution of in-water structure.

Fortunately, we know from Figure 2 that we can remove most of the striping and speckling artifacts by simply fixing the aerosol model. This suggests that a good result can be obtained if we can stabilize the relative variability between the two NIR bands, leaving the 865 nm channel unchanged. Ideally, we would like to minimize the small-scale variability in $\varepsilon(765,865)$, the ratio of single-scattering aerosol reflectance between 765 and 865 nm, as this is the ultimate index into the aerosol model selection process. We can't do that directly within MSL12, due to the complexities of determining the single-scattering $\varepsilon(765,865)$, but we can minimize the small-scale variability in multi-scattering epsilon, $\varepsilon_{ms}(765,865)$.

The proposed smoothing algorithm is as follows:

1) define NIR aerosol radiance at pixel i for wavelength λ as:

$$L_a(\lambda,i) = [\ (L_t(\lambda,i) - tL_f(\lambda,i))/t_{O3}(\lambda,i) - L_r(\lambda,i)\ ]/t_{O2}(\lambda,i)$$

2) given a scan/pixel window centered on pixel x, containing a total of n unmasked pixels, compute mean $L_a(\lambda)$ at x as:

$$<L_a(\lambda,x)> = 1/n\ \ \Sigma\ L_a(\lambda,i),\ \ \text{for } i=1,n \text{ and } \lambda=765 \text{ or } 865\text{nm}$$

3) compute mean, multi-scattering epsilon at pixel x as:

$$\varepsilon_{ms} = <L_a(765,x)> / <L_a(865,x)>$$

4) now compute a new $L_a(765,x)$ which would yield the mean epsilon when combined with the original $L_a(865,x)$:

$$L_a'(765,x) = \varepsilon_{ms}\ L_a(865,x)$$

5) and finally, reconstruct the TOA radiance at 765nm:

$$L_t(765,x) = [L_a'(765,x)\ t_{O2}(765,x) + L_r(765,x)]\ t_{O3}(765,x) + tL_f(765,x)$$

where:

$L_t(\lambda,i)$ is the observed TOA radiance for wavelength λ at location i
$tL_f(\lambda,i)$ is the white-cap radiance, transmitted to the TOA
$L_r(\lambda,i)$ is the Rayleigh path radiance
$L_a(\lambda,i)$ is the aerosol path radiance, including Rayleigh-aerosol interaction
$t_{O2}(\lambda,i)$ is the oxygen transmittance
$t_{O3}(\lambda,i)$ is the ozone transmittance

The algorithm described above has been implemented as a filter function in MSL12. The effect is to replace each TOA radiance observation at 765 nm with a new value which minimizes the spatial variability in ε_{ms}, the multi-scattering equivalent of atmospheric epsilon in the NIR. The new TOA radiance spectra can then be processed in the standard manner through the atmospheric correction algorithm, with the expectation that the spatial variability in atmospheric model selection should be significantly reduced. Some results of applying a 5x5 square epsilon-targeted smoothing filter on our test subscene are provided in Figures 5 and 6.

Comparing the Ångstrom(510,865) image of Figure 5 with the original Ångstrom retrieval of Figure 1, we see a dramatic reduction in pixel-to-pixel variability. As expected, the spurious high and low-valued Ångstrom retrievals have been effectively eliminated, while the large-scale and meso-scale variability, such as the change of aerosol type toward the coastal areas and in localized regions, has been maintained. The speckling and striping effects in the nLw(443) retrievals have also been substantially reduced, as can be seen by comparing Figure 6 with Figure 2. In fact, the spatial structure of Figure 6 looks very much like that of Figure 3, the fixed aerosol model case.

Figure 5: *Ångstrom Coefficient at 510 nm, 5x5 epsilon mean filtering* .

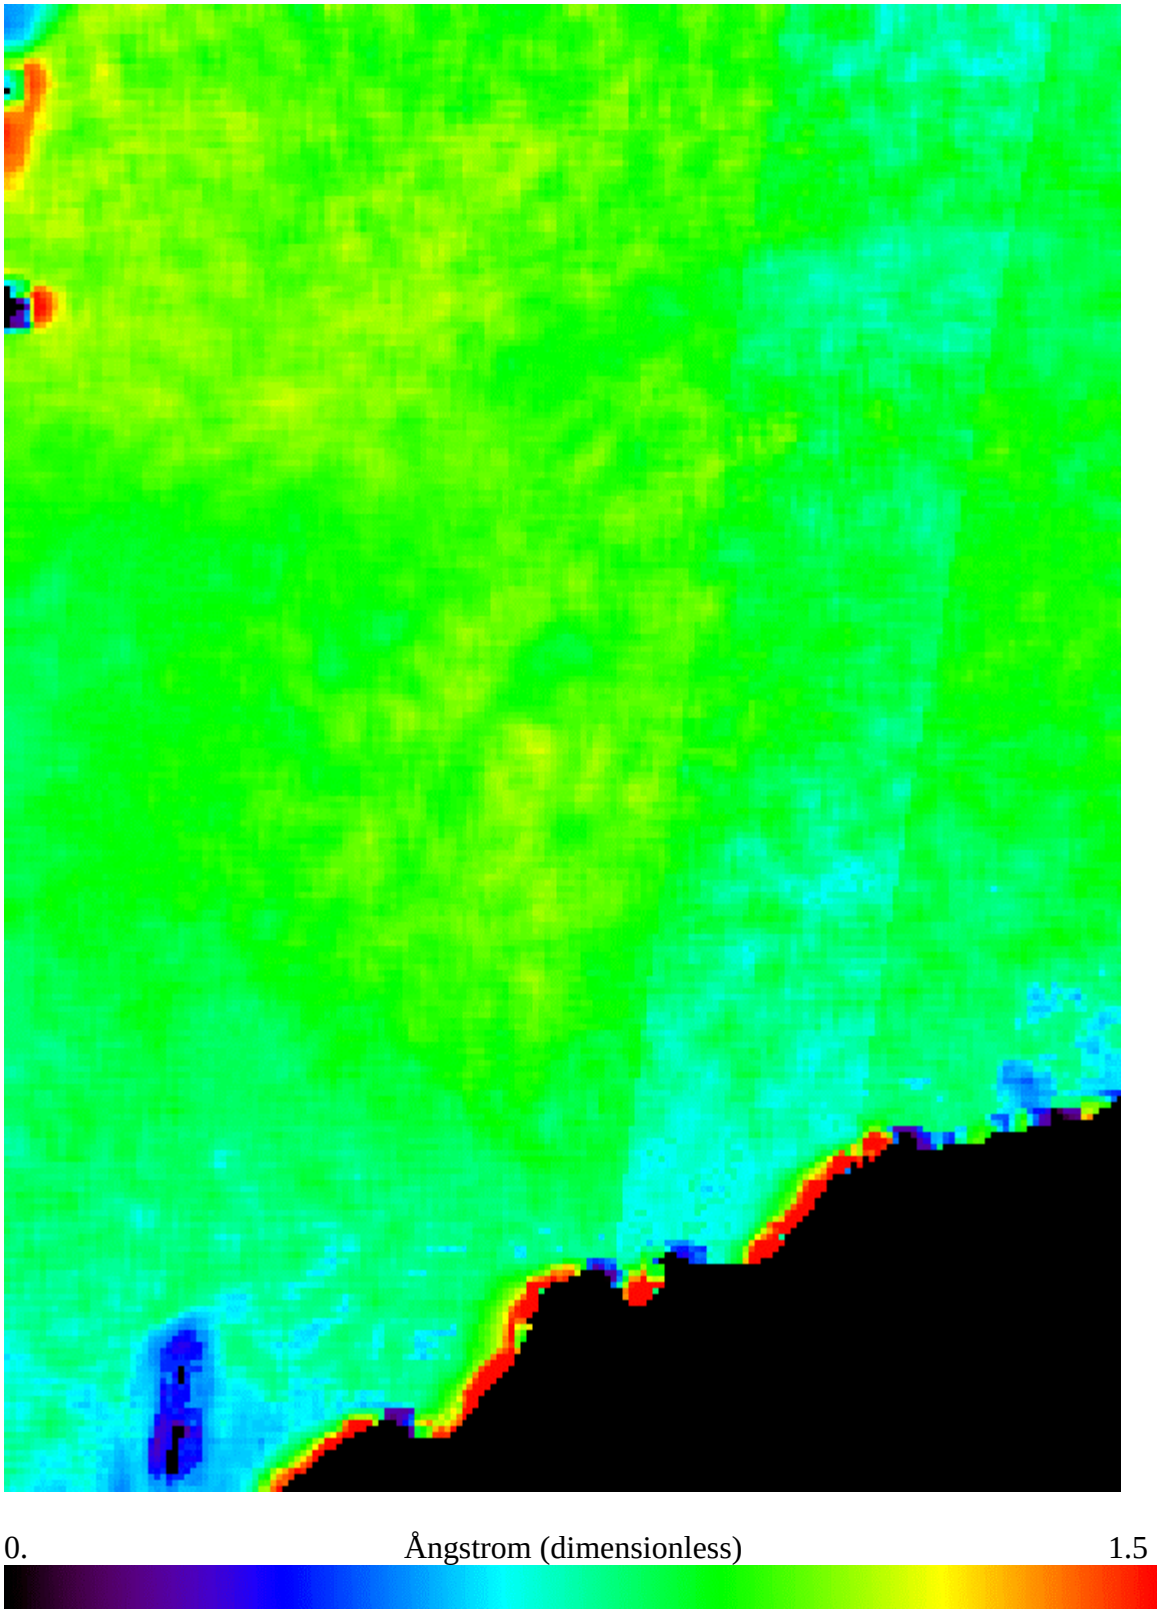

Figure 6: *Normalized water-leaving radiance at 443 nm, 5x5 epsilon mean filtering .*

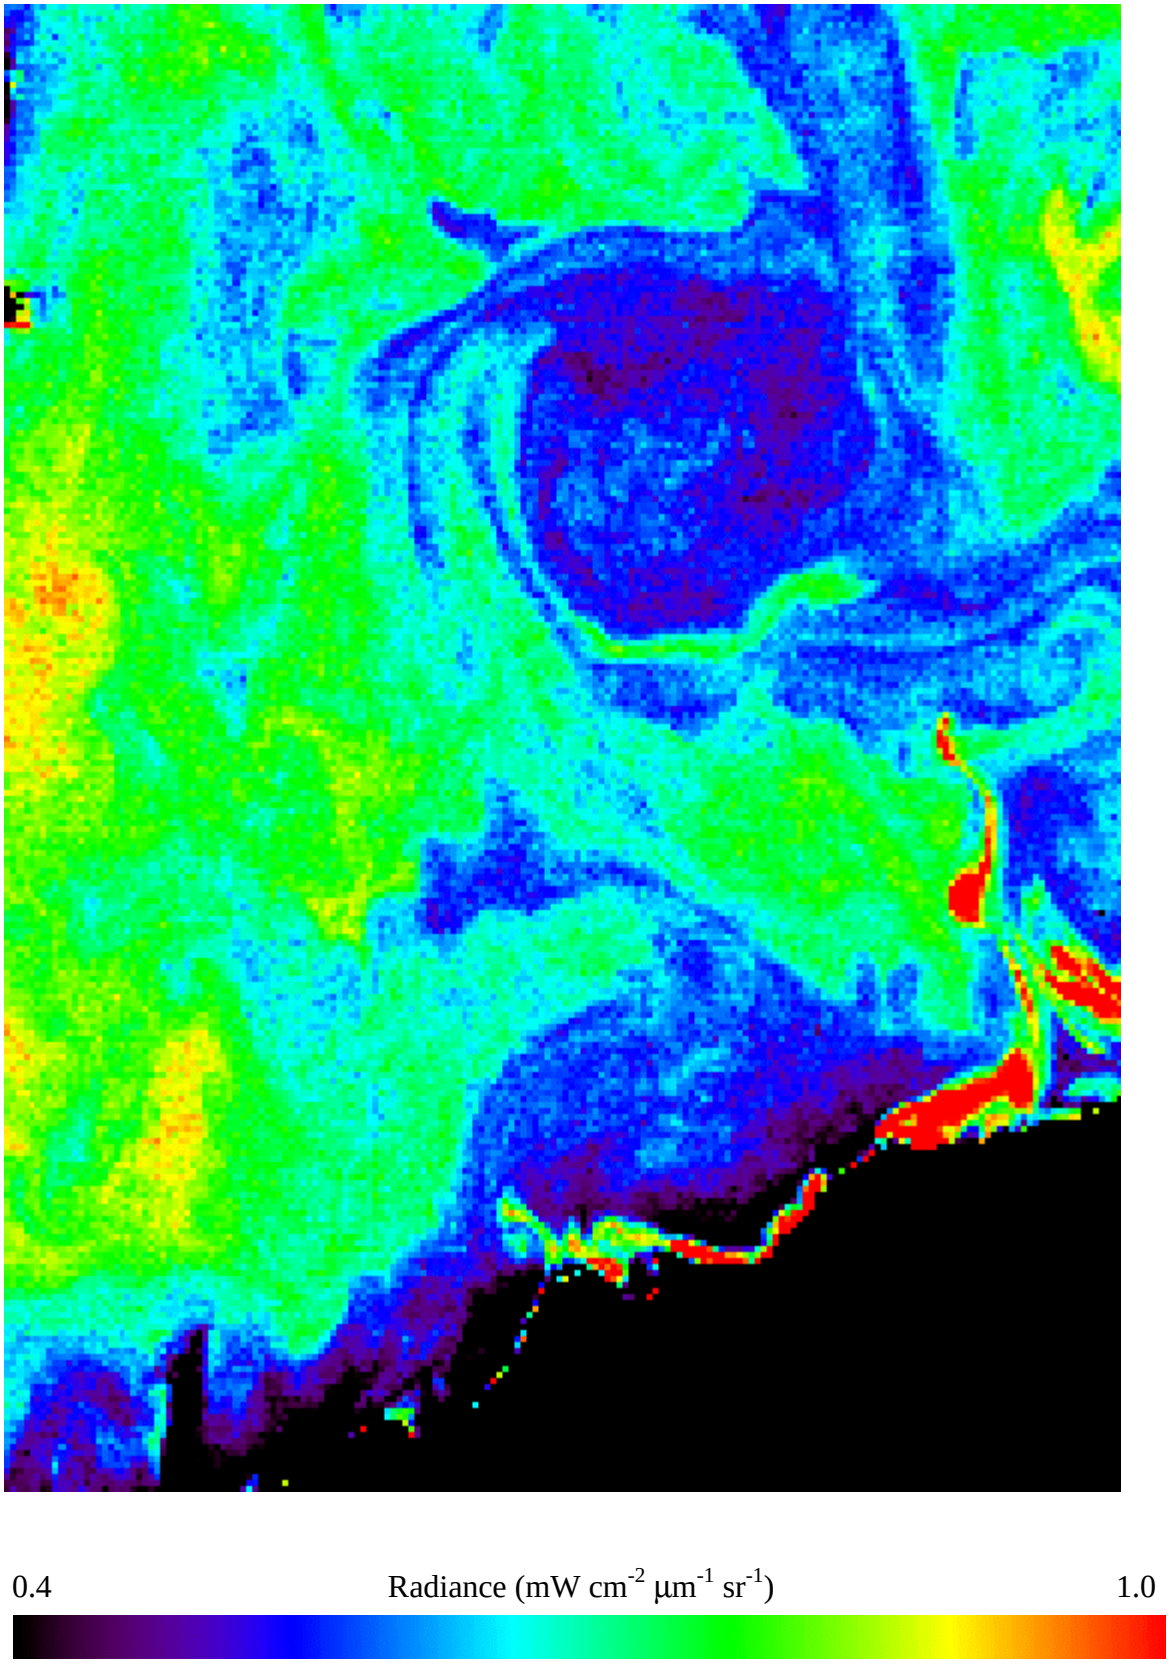

IV. Additional Refinement

The above results were derived using a 5-pixel by 5-line square window to compute the filtered value of $L_t(765)$. It is desirable to keep the window size as small as possible, to minimize the reduction of real changes in aerosol type. However, the window size obviously needs to be large enough to allow sufficient sample size for the averaging to be effective. Additional studies have found that the 3x3 square window is not sufficient (http://orca.gsfc.nasa.gov/pub/seawifs/smooth/filtering.html). An alternative to changing the size of the filter is to change the shape, from square to diamond. This can be achieved by introducing the concept of a filter window kernel, which indicates which pixels within the window will be considered in computing the filtered value. Consider these two examples of a 5x5 filtering window, where the value of 1 indicates that the pixel at that location will contribute.

Square 5x5				
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1

Diamond 5x5				
0	0	1	0	0
0	1	1	1	0
1	1	1	1	1
0	1	1	1	0
0	0	1	0	0

For the same window size, the diamond filter kernel reduces that number of contributing samples by approximately 50% over the square kernel, and the radius of influence is never greater than two pixels. While this obviously reduces the number of samples contributing to the mean, the diamond shape is better designed to minimize the specific, line-by-line digitization problem that is seen in the SeaWiFS data. This is because the diamond kernel gives nearly equal weight to the odd and even lines, while the square kernel yields a 3 to 2 over-weighting of opposing lines. Additional discussion of the relative merits of the diamond filter kernel can be found here: http://orca.gsfc.nasa.gov/pub/octs/gac/octs_filtering.pdf.

Figures 7 and 8 show the effect of the 5x5 diamond epsilon mean filtering on Ångstrom(510,865) and nLw(443). Careful comparison of these two images with those of Figures 5 and 6 will show that the diamond filter retains mesoscale structure in the Ångstrom field, while better reducing unnatural line-by-line structure in the nLw(443) image.

Figure 7: *Ångstrom Coefficient at 510 nm, 5x5 diamond epsilon mean filtering* .

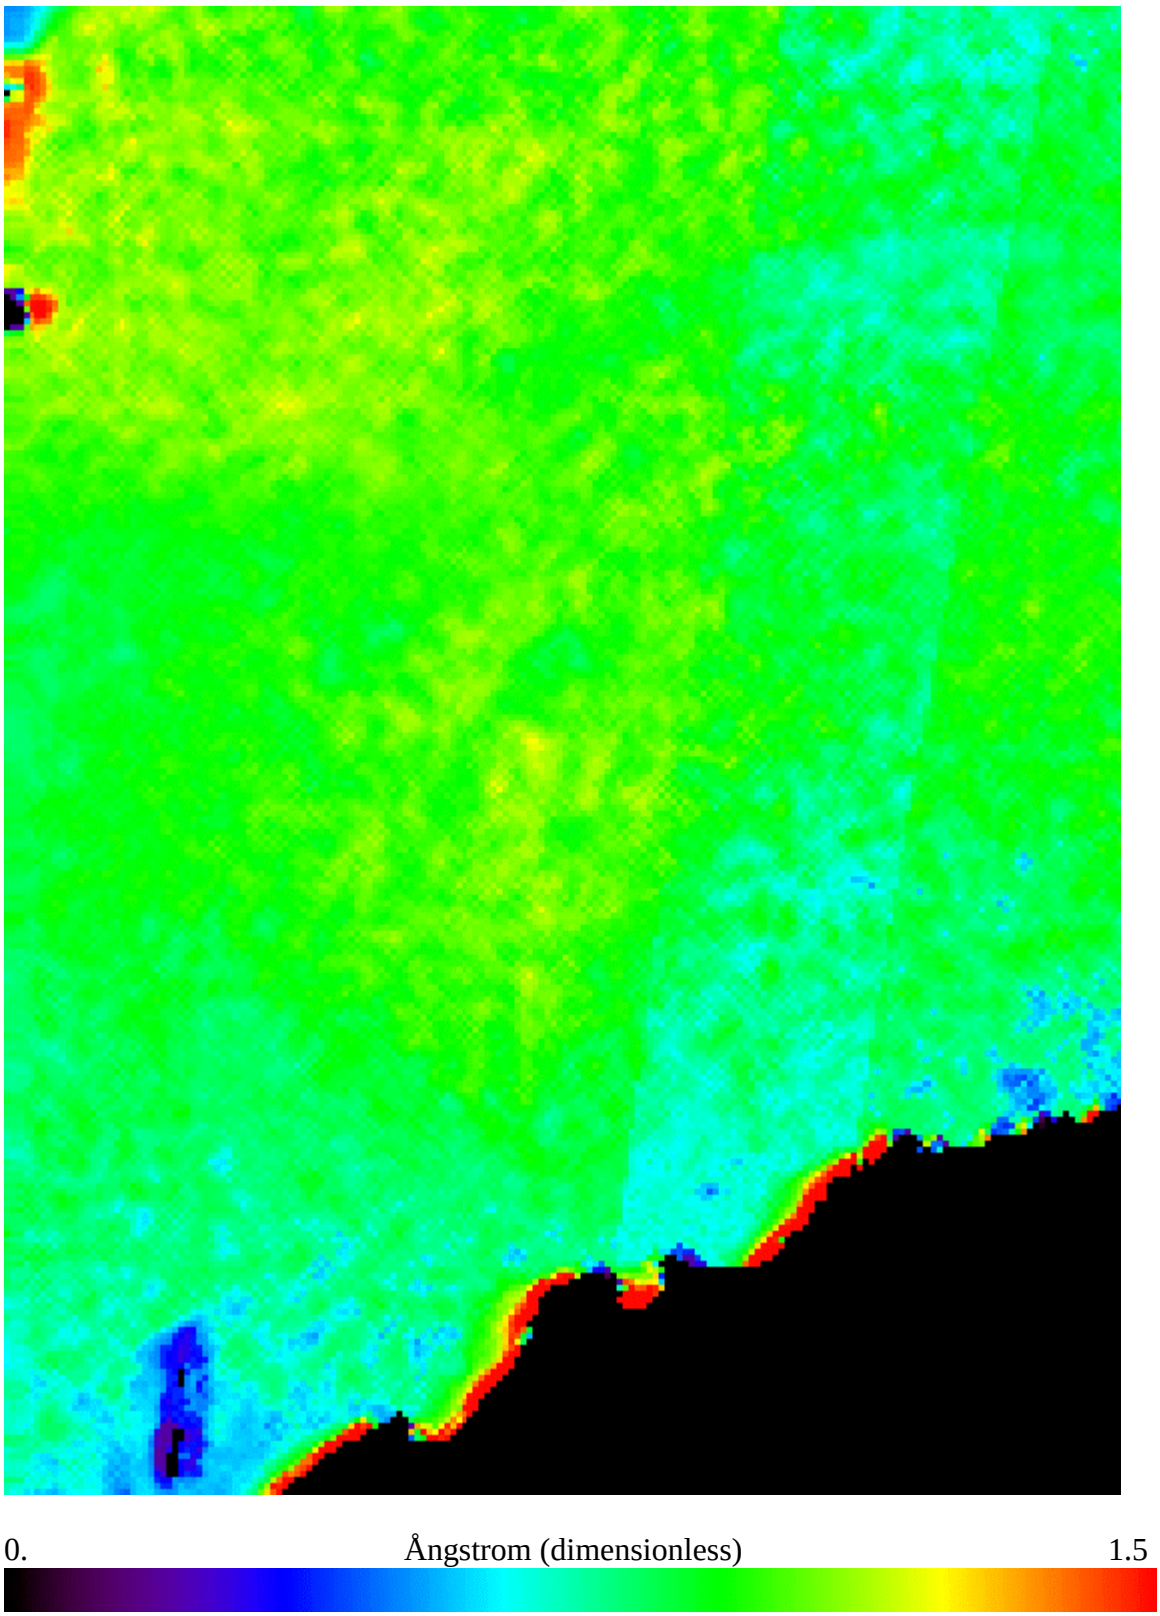

Figure 8: *Normalized water-leaving radiance at 443 nm, 5x5 diamond epsilon mean filtering*

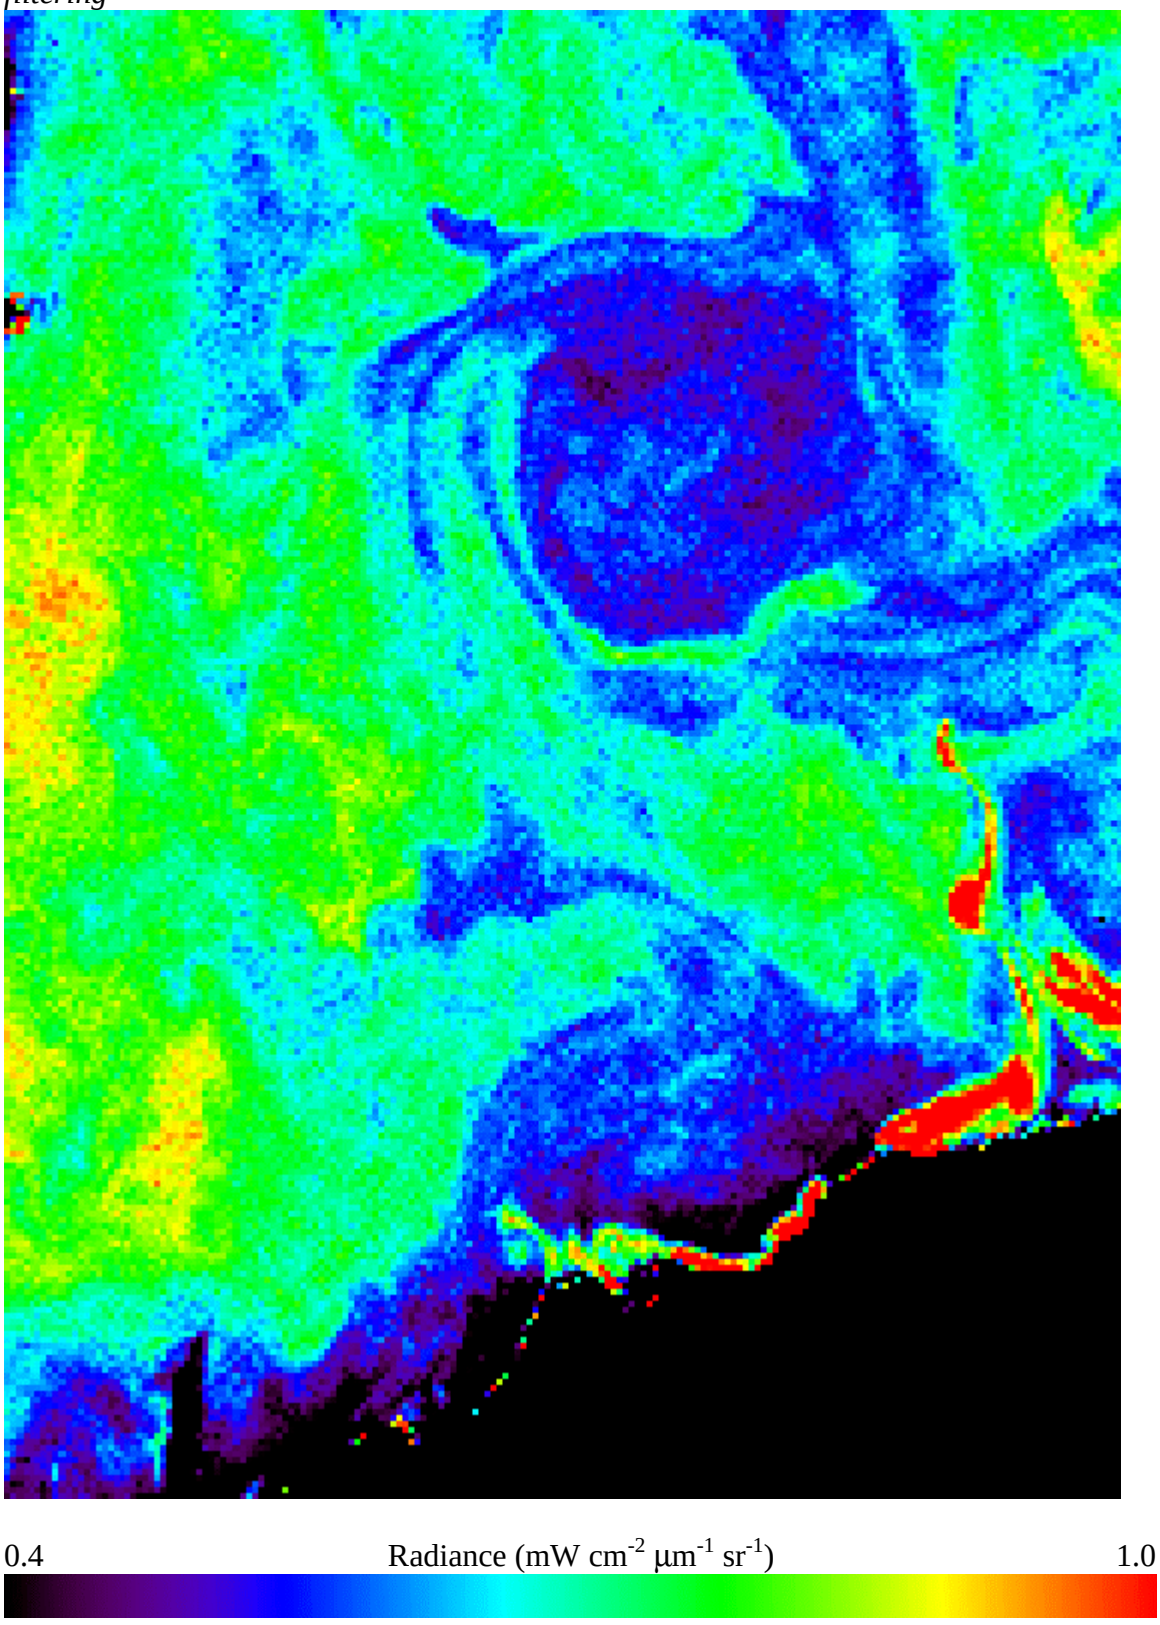

Conclusions